Niche specificity, polygeny, and pleiotropy in herbivorous insects

Nate B Hardy

Department of Entomology and Plant Pathology, Auburn University

Matt Forister

Department of Biology, University of Nevada, Reno

Abstract

What causes host-use specificity in herbivorous insects? Population genetic models predict specialization when habitat preference can evolve and there is antagonistic pleiotropy at a performance-affecting locus. But empirically for herbivorous insects, host-use performance is governed by many genetic loci, and antagonistic pleiotropy seems to be rare. Here, we use individual-based quantitative genetic simulation models to investigate the role of pleiotropy in the evolution of sympatric host-use specialization when performance and preference are quantitative traits. We look first at pleiotropies affecting only host-use performance. We find that when the host environment changes slowly the evolution of host-use specialization requires levels of antagonistic pleiotropy much higher than what has been observed in nature. On the other hand, with rapid environmental change or pronounced asymmetries in productivity across host species, the evolution of host-use specialization readily occurs without pleiotropy. When pleiotropies affect preference as well as performance, even with slow environmental change and host species of equal productivity, we observe fluctuations in host-use breadth, with mean specificity increasing with the pervasiveness of antagonistic pleiotropy. So, our simulations show that pleiotropy is not necessary for specialization, although it can be sufficient, provided it is extensive or multifarious.

Introduction

Most species of herbivorous insects use as hosts only a small fraction of the plant species they encounter in their environment (Forister et al., 2015; Futuyma & Moreno, 1988), despite the obvious advantages of using many hosts and having a broader resource base (Hardy et al., 2020; Rainey et al., 2000; Smith & Hoekstra, 1980). For the most part, theorists have explained the evolution of this specificity as a consequence of antagonistic pleiotropy (Egas et al., 2004; Rueffler et al., 2004): an allele that confers high performance on one host confers low performance on another. Mathematical models show that if host-use performance is determined by one gene with antagonistic pleiotropy, and if populations can evolve host-use preferences, then the evolution of host specialization is all but inevitable, at least for panmictic populations in static environments (Ravigné et al., 2009).

But in two key aspects, this theory bears little resemblance to what we know about herbivorous insects. First, in herbivorous insects, most host-specific performance phenotypes, such as the ability to develop on a diet of any particular plant species, seem to be controlled by many genes, not by one or even a few (Egan et al., 2015; Vertacnik and Linnen 2017). Second, empirical studies of polygenic local adaptation (reviewed by Savolainen et al. 2013) show that antagonistic pleiotropy tends to be rare in comparison to conditional neutrality; alleles tend to affect performance in one environment but not another. As a striking example, in the Melissa blue butterfly, *Lycaeides melissa*, more than 99% of alleles with host-dependent performance effects are conditionally neutral (Gompert et al. 2015). Although linkage disequilibrium can also cause negative genetic covariances (Via & Lande, 1985), conditionally neutral alleles might seem less likely than pleiotropic alleles to drive the evolution of host-use specialization. Moreover, it is not obvious what influence a few – or many – pleiotropic loci might have in the evolution of a multi-dimensional quantitative trait such as host-use. Here, we use forward-time individual-based quantitative genetic simulations to address this simple question: do

simple genetic trade-offs matter in a polygenic world? More specifically, how might pleiotropy affect the evolution of specialization, when preferences can evolve along with performance?

Our models build on much previous work. The theory of the evolution of niche specialization could fill volumes. Ravigné et al. (2009) classify more than seventy studies by five factors: (1) the specific research question addressed, (2) the type of population regulation, (3) the shape of the local-adaptation trade-off, (4) whether or not habitat choice was allowed to evolve, and (5) if so, what the mechanism was for the evolution of choice. Our question is how the genetic architectures of habitat preference and performance affect the evolution of niche specialization. In that respect, our work departs from what has come before, almost all of which has focused on the emergence or maintenance of local-adaptation polymorphisms with single-locus architectures for performance and/or preference. Of the few previous studies that take a quantitative genetic approach (e.g. Via and Lande 1985; van Tienderen 1991, 1997), none examine the joint evolution of habitat preference and performance. Also, because our models are individual-based, population regulation and adaptive trade-offs are not as they are in population-based models. Concretely, in our models, population size is something that emerges via the balance of individual mortality, viability, and fecundity. Likewise, adaptive trade-offs – that is, quantitative genetic covariances – are not fixed, but rather emerge via genetic linkage and the pleiotropic effects of specific mutations, and evolve via subsequent mutation, recombination, genetic drift, and natural selection. As for the mechanisms of habitat choice, our models resemble so-called "two-allele" models (Felsenstein 1981), although in this case many more than two alleles are involved. Habitat choice evolves at least somewhat independently from local adaptation, and preference alleles affect proclivity for a specific environment, rather than attachment to wherever they happen to have come from, or whatever would be ideal for their performance.

Although we are primarily interested in the evolution of niche breadth, since our models allow for ecologically-driven population genetic divergence without geographic isolation, they also speak to

sympatric speciation, the theory of which is based on many of the same assumptions as that of niche specialization, and was reviewed by Kirkpatrick and Ravigne (2002) and Bolnick and Fitzpatrick (2007). Previous sympatric speciation models vary by (1) the causes of disruptive selection, (2) how individuals select mates, (3) whether selection acts directly or indirectly on mating traits, and (4) the genetic basis of mating traits. In our model, disruptive selection stems from resource heterogeneity, mating is assortative, and selection acts indirectly on the mating trait, namely host-preference (mating takes place on the host), as fitness is only affected by a host-dependent performance phenotype. To our knowledge, our model is the first to have this particular combination of features, or to allow for conditional neutrality at quantitative trait loci.

Our models show that for pleiotropy to drive quantitative specialization, it needs to be quite extensive. Therefore, we also consider what else might drive specialization in sympatry. For direction, we turn to the foundational quantitative genetics work of Via and Lande (1985). With an analytical approach, they identified three conditions for specialization, to wit, there is (1) no gene flow between habitats, (2) insufficient genetic diversity for the relevant niche traits, or (3) a genetic correlation across habitats equal to ± 1, which would be true if phenotypes are monogenic. Otherwise, populations evolve high-performance generalism. Nevertheless, the *rate* of adaptation can vary across habitats. For example, Via and Lande (1985) show that if habitat types vary in frequency, adaptation will be faster and more direct in the more common habitat, something which has also been demonstrated with phenotype-free population genetic models (Whitlock 1996; Kawecki et al. 1997). Moreover, until the population evolves a near-optimal phenotype in the common habitat, the evolutionary response in the rare habitat is largely correlational. And if genetic covariances are negative, this can temporarily increase maladaptation to the rare habitat. To sum up: barring monogeny, monophorphism, and allopatry, given enough time, habitat generalism will evolve, but along the way negative genetic

covariances and asymmetries in productivity between habitats increase the odds of temporary specificity.

These insights are contingent on several model assumptions: that (1) habitat-specific phenotypic optima are fixed, (2) habitat preference does not evolve, (3) selection is weak and (4) population size is infinite, so that (5) the matrix of genetic variances and covariances is fixed. Our models relax each of assumptions. Therefore, they provide a more general – although less elegant – view of quantitative niche breadth evolution.

Methods

Simulations were performed with SliM 3 (Haller & Messer, 2019). In this framework, models are specified in the Eidos programming language. Eidos codes are available from Zenodo (https://doi.org/10.5281/zenodo.6819401). For a list of model parameters and variables see table 1.

Initialization (fig1a). At the start of each simulation, we initialize a population composed of n=2 subpopulations, each of which is composed of N_j hermaphroditic individuals. We imagine that these individuals are insect herbivores, and that each subpopulation is associated with host species h_j. We give each individual the same diploid 100 kb, one-chromosome genome. This initial genome is an empty container for future mutations. It is divided into ten 5-kb coding regions interspersed with ten 5-kb non-coding regions. Host-preference genes are mixed in with host-performance genes in the coding regions of the genome, and in some models the same loci affect both traits. Of course, real herbivorous insect genomes are much larger and have multiple chromosomes. Working with smaller genomes is easier computationally, but increases linkage disequilibrium. To correct for that, we ran some simulations with a recombination rate, χ, of 1e-4 events per base pair (bp) per generation. At that rate, we expect about ten crossing-over events per genome per generation, which corresponds roughly to as much genetic shuffling as would be typical for a real herbivorous insect species (Wilfert et al., 2007).

This assumes that preference and performance alleles are randomly distributed across the genome. To assess the impact of tighter linkage, we also ran simulations with χ=1e-6.

In most simulations, the mutation rate, μ, is 1e-7 per bp per generation, yielding an expected 0.01 non-neutral mutations per individual per generation. This is an order of magnitude more than what we would expect based on empirical measures of mutation rates (~6e-9 per bp per generation; Liu et al. 2016) and average gene size (1,346 bp; Xu et al. 2006), and if we assume that quantitative host-use performance and preferences are each controlled by ~100 genes (Vertacnik and Linnen 2017). But an elevated per-individual mutation rate compensates for the fact that our simulated populations are smaller than most real ones. Four types of mutations can occur: (1) neutral mutations in non-coding regions, (2) pleiotropic mutations, that is, mutations at quantitative trait loci (QTLs) that affect performance on each host species, and in some models also affect host preference, (3) conditionally neutral mutations at QTLs that affect performance only on host h_j, and (4) conditionally neutral mutations at QTLs that affect preference for host h_j.

After initialization, we simulate many generations following the life cycle shown in fig. 1.

1. Mating. At the start of each generation, each surviving individual randomly selects as a mate one individual from their subpopulation. This reflects the tendency of herbivorous insects to mate on or near their host plants. The fecundity of the choosing individual is determined by a draw from a Poisson distribution with a mean and variance (λ) of 1.5 (fig. 1b). The drawn number of crosses between a mating pair (chooser and chosen) are then added to the subpopulation. Note that choosers can also be chosen as mates by other individuals in their subpopulation; in theory, an individual could be randomly chosen as the mate of every other individual in their subpopulation, but the chances of being involved in more than two or three matings is small. As is typical for herbivorous insects, generations are discrete and non-overlapping; after mating and producing offspring, parents die.

2. Offspring production. This entails mutation and recombination of parental genomes. When a mutation occurs at a conditionally neutral performance or preference QTL, an allele effect is drawn from a univariate normal distribution with a mean of zero and a variance (σ^2) of 0.6 (fig. 1c). When a mutation occurs at a pleiotropic performance QTL, its effects are drawn from an bivariate normal distribution (fig. 1d) with means of zero, variances of 0.6, and symmetrical covariances, κ, which varied over simulations $\{-0.9\sigma^2, 0, 0.9\sigma^2\}$. With these distributions of allele effects, on average, simulated populations adapt to changing host environments via promotion of on the order of a few tens of alleles. As previously mentioned, we also examined models in which pleiotropic mutations affected host preference as well as performance in each host environment. For these simulations pleiotropic allele effects are drawn from a tri-variate normal distribution with means of zero, variances of 0.6, the same range of performance trait covariances as for the two-dimensional models, and zero covariances between the preference and performance traits. Our main goal was to see how host-use evolution depends on the ratio of pleiotropic to conditionally-neutral alleles governing host-use performance, and (in some cases) also preference. Therefore, across simulations, we varied this mutational mix, M, $\{0, 2^{-3}, 2^0, 2^3, 2^6\}$.

The relative rates of preference and performance mutations, ρ, could affect the odds of population divergence and specialization. Indeed, if mutations that affect host preferences are common, they could be an adaptive hazard, that is, a genetic architectural liability that increases the risk of maladaptation (Lynch 2011). For models with performance-only pleiotropies, we considered two values for ρ spanning a range of plausibility $\{0.1, 1\}$; this was based on our reading of the relevant gene expression research literature (e.g., Andersson et al. 2014; Wu et al. 2016; Yu et al. 2016; Tan et al. 2019). To get a sense for the sensitivity of some model results to the effects of performance-affecting alleles, we also ran some simulations with ρ=0. For models with performance-and-preference

pleiotropies, we considered only a 1:1 ratio of conditionally-neutral mutations affecting preference and performance.

3. Migration and habitat choice. Each newly-produced individual selects a host species. At the start of a simulation, all individuals have a preference genotype value, y_1, of zero, and the probability of choosing h_1 is 0.5. Then, as the simulation progresses, host preference evolves as a function of the sum of allele effects at preference-determining loci (fig. 1e). Concretely, the probability that an individual will choose h_1 over h_2 is given by a logistic function, $P(y_1) = - 1 / (1 + \exp(-1 * y_1)$. This allows from some geneflow as mean preference phenotypes slowly approach either of the boundary states, $0 \leq P(y_1) \leq 1$.

4. Selection and regulation. The next step in the life cycle is hard selection. An individual's viability, that is, their probability of surviving long enough to mate, is determined by two factors: (1) the match between their performance genotype and host-environment (fig. 1f), (2) local subpopulation density. An individual's performance genotype for host h_j is the sum of allele effects across pleiotropic loci and the conditionally neutral loci affecting performance only that host. We assume no epistasis or phenotypic plasticity. Performance phenotypes map to viability with a standard Gaussian fitness function:

For $X \sim N(Oj, \omega)$, $v(z_{ij})=P(z_{ij})/P(Oj)$,

where z_{ij} is an individual's genotype value for performance on h_j, $v(z_{ij})$ is the viability of an individual with z_{ij} on h_j, O_j is the optimal value for the performance phenotype on h_j, ω is the variance {3.0}, that is, the steepness of the selection gradient (smaller variances map to steeper gradients), and division by $P(O_j)$ sets the viability of a perfect match to one. The force of selection on a population depends on both ω and the rate, δ_O, at which host optima move away from their initial values {0.005, 0.05, 0.1}.

After genotype-environment matching, individual viability is re-scaled by local subpopulation density. Specifically, each host h_j has a nominal carrying capacity, K_j, which is set to 500 individuals.

Viabilities are re-scaled by the ratio K_j to the current subpopulation size N_j. By capping this density scaling factor at 1.0, we make selection hard; when $N_j > K_j$, the viability of all individuals on a host species is reduced, but low density ($N_j < K_j$) is not enough to drive demographic growth of maladapted genotypes. So, N has a soft upper limit equal to the sum of the nominal carrying capacities of host species. If selection is strong enough that the product of mean viability and fecundity is < 1, the population will shrink, and if that product remains < 1, the population will go extinct.

After selection, the life cycle again starts again.

The arc of a simulation. At the start of each simulation the optimal value for each performance phenotype O_j, is set to zero, which matches the starting genotype value of each individual in the population. Thus, herbivore populations start off well-adapted to their host environments. O_j stays at zero for a 1,000 generation burn-in period in which the simulated population accumulates genetic diversity. Then, in each of 1,000 subsequent generations, O_j shifts by a random draw from an exponential distribution with mean δ_O. Such gradual and directional changes approximate a variety of real-world situations, for example the effects of global warming, land-use intensification, or biological invasion, and allow host-use adaptation to proceed via incremental changes to mean genotype values.

Model variants and statistics. We examined two main models. In Model 1, pleiotropic mutations affect only performance. In Model 2, pleiotropic mutations affect both host-use performance and preference. For both models, we explore various combinations of values for δ_O, M, ρ, χ, and κ (Table 1). For Model 1, we also examined the effects of asymmetries in the productivity of host environments; fixing the global carrying capacity to 1,000, and letting the proportional carrying capacity of h_1, pK_1 vary {0.05, 0.25, 0.5}.

For different perspectives on the genetic states of the population at the end of each simulation under Model 1, we recorded the mean values over a simulation for these statistics: (1) N, total population size across hosts, (2) H, heterozygosity, (3) H', host-use diversity of the herbivore

population, characterized with Shannon's diversity index, (4) W: mean population fitness across host-associated subpopulations; (5) (5) F_{ST}: the standard ratio of between-subpopulation neutral genetic variance to total neutral genetic variance; (6) δP, the difference in mean preference phenotypes between host species, (7) δW, a specialization index that measures subpopulation divergence, defined as the mean difference between an individual's fitness on their host species and what their would have been on other host species, scaled by the overall mean fitness of the population (ranging from zero to one, with smaller values denoting more generalist populations). Below, we focus on whichever statistics we found most revealing.

Before we look at any results, here is some guidance on how to interpret the pleiotropic effect covariances. After the burn-in period, the optimal value for the host-performance phenotype for each host species, O_i, increases by δ_O each generation, and the mean performance trait values in the herbivore population tend to lag behind these diverging optima. Thus, pleiotropic performance mutations drawn from a bi-variate normal distribution with positive covariances tend to move the phenotype either towards or away from the diverging host-use optima (fig. S2), in other words, these are "general vigor" or "general weakness" alleles (Fry 1996). In the former case, pleiotropy can help generalists solve the problem of adapting to a heterogeneous host environment. In contrast, with negative pleiotropic covariances, alleles that increase performance on one host tend to decreases it on others, that is, antagonism abounds. With covariances of zero, positive and antagonistic pleiotropy are equally likely. For Model 2, the covariances between preference and performance trait effects are zero; in other words, alleles that affect performance also affect preference, but positive and negative relations are equally likely between effects on preference and the effects on each performance phenotype.

Results and Discussion

Model 1. Performance pleiotropies. Let us first consider what happens when the changes in host-performance optima are slow (δ_O=0.005). Figs. 2a-d show an example run, where there is an even

mix of pleiotropic and conditionally neutral mutation (M=1), pleiotropic covariances are zero (κ=0), the recombination rate is what is expected with a uniform distribution of QTLs across the genome (χ=1e-4), and performance and preference mutations occur at an equal rate (ρ=1). Population genetic summary statistics vary from one generations to the next, but without strong directional trends; δW stays near zero (no specialization), W stays near unity (high fitness), N hovers around $2K_j$, and H' stays near its maximum value. The population maintains high-performance generalism. Figs. 3a-d summarize batches of 20 simulations for each combination values for M, and κ given in Table 1, and with ρ=1.0. Each plot shows how a population variable tends to change with increasing amounts of pleiotropy in the mutational mix (M), as well as how those changes depend on the covariances of pleiotropic effects (κ). Unless antagonistic pleiotropy is many times more common than conditional neutrality (M>>1), populations maintain high-performance generalism. W is near unity, H' is near its maximum value, and δW is near zero. Subpopulations do divergence some in host preference, with $\delta P \sim 0.2$, but this is firmly on the non-choosy end of the spectrum. Even with pervasive and antagonistic pleiotropy (M=2^6; κ=-$0.9\sigma^2$) the evolution of host-use specialists is unlikely. Moreover, high levels of antagonistic pleiotropy tend to reduce mean fitness; thus insofar as mutational covariances are evolvable, selection will be against antagonism. Antagonistic pleiotropy is even less likely to drive specialization or strong host preferences when preference mutations are less common than performance mutations, ρ=0.1, (figs. S2a-d). Lowering the recombination rate between performance alleles (χ=1e-6) increases the scope for specialization via pleiotropy, but even so, only when ρ=1.0, M>1 and κ=-$0.9\sigma^2$ (figs. S3a-d). In summary, with slow changes in performance optima, barring exceptionally tight linkage, negative genetic covariances, and high preference mutation pressure, quantitative specialization via any amount of pleiotropy is unlikely.

Now consider what happens when performance optima diverge rapidly. Figs. 2e-h show an example run in which M=0 (there is no pleiotropy), δ_O=0.1, and ρ=1. See the population veer away

from the joint optimal performance vector; see the rise of δW and the fall of N and H'. We stopped the simulation at generation 1206, at which point the $N_1 = 0$, and the population is composed exclusively of h_2 specialists. Figs. 3e-h**,** summarize a batch of runs in which we vary M and κ while $\delta_O = 0.1$, χ=1e-4, and ρ=1.0. Even when pleiotropy is rare and neutral, populations tend to evolve some level of host-use specialization; $\delta W \sim 0.5$, and H' is about half of its maximum value. The effect of M depends on κ; increasing M increases δW and δP when κ=- $0.9\sigma^2$ and decreases it when $\kappa = 0$ or $0.9\sigma^2$. Also note that regardless of the value of κ, increasing the amount of pleiotropy in the mutational mix (M) increases mean viability (W); in one way or another, depending on its sign, pleiotropy helps the metapopulation adaptively keep pace with the changing host environment. On the one hand, positive pleiotropies provide a double-shot of the adaptive genetic diversity that is depleted by strong and sustained directional selection. On the other hand, antagonistic pleiotropies can push for the evolution of strong host preference, which increase rates of host-use adaptation by effectively homogenizing the host environment to which a genotype will likely be exposed.

Decreasing χ further increases the odds of specialization with little pleiotropy (figs. S3e-h), whereas decreasing ρ does the opposite (fig. S2e-h). In summary, when the host environment changes rapidly, the odds of host-use specialization increase across genetic architectures. But antagonistic pleiotropy is not necessary. Specialization can also evolve via tight linkage of conditionally neutral alleles, or via high preference mutation pressure.

Another thing that increases the odds of specialization is asymmetry in the productivity of host environments. Figs. 2i-l show an example run for the case in which $pK_1 = 0.05$ – that is K_1 accounts for 5% of the total carrying capacity across hosts – and where M=0, ρ=1, δ_O=0.005, and χ=1e-4. As per the classical theory of Via and Lande (1985), performance evolution is initially dominated by adaptation to the more common host, h_2. But with the evolution of divergent specialization between subpopulations, the population's evolutionary path bends towards the joint optimum vector. In this particular

simulation, the population passes through this phase of specialization and begins to coalesce and re-generalize. In other runs (not shown), subpopulations stay specialized; the population's evolutionary path runs parallel to and above (with z_2 on the y-axis) the joint optimal performance vector for the duration of the 1,000 generation observation period. But the fact that populations do occasionally coalesce resonates with previous work demonstrating the importance of gene flow from source subpopulations for adaptation in marginal sink subpopulations (Holt et al. 2003).

Figs. 3i-l summarize a batch of runs with pK_1 = 0.25, ρ =1, χ = 1e-4, and δ_O =0.005, varying M and κ. Here again we see that even without pleiotropy populations can evolve intermediate levels of host-use specificity (δW ~ 0.1), although to a lesser extent then when δ_O=0.1 or χ=1e-6. And here again, it is not until M > 1 that antagonistic pleiotropy really begins to affect specificity, in this case causing a precipitous decline in H'. When host species vary in productivity, pervasive antagonistic pleiotropy tends to push for specialization on the more productive host.

Model 2. Preference-and-performance pleiotropies. Fig 4 shows an example run with χ = 1e-4, ρ = 1, κ = - $0.9\sigma^2/2$, $M = 2^{-3}$, and δ_O = 0.005. Note that there are stretches in which both host environments are occupied by non-choosy, well-mixed, and well-adapted generalists. And there are stretches in which the population evolves strong preference for one host species or the other, and performs well on that host species, and the other host species is nearly or completely unused.

We summarize our analysis of Model 2 with the plots in fig. 5. Here χ = 1e-4, and ρ = 1. In 5a-d, μ = 1e-7 and δ_O = 0, that is, the optimal performance phenotype is the same in both host environments and does not change over the simulation. Without pleiotropy (M = 0) – in which case all preference mutations are conditionally neutral – simulated populations maintain high-performance generalism. With increasing preference-and-performance pleiotropy, populations spend more time as specialists on one host species or the other; H' declines steeply with increasing values for M, regardless of κ. (δW stays low, but given that performance optima are identical for each host species, this is as

expected). Now look at figs. 5e-h, where $\delta_O = 0.005$ and μ=1e-9, that is, when the host-performance optima diverge, and we drop the rates of all kinds of mutation by two orders of magnitude. In this case, populations tend to have generalized host-use regardless of pleiotropy; even when $\kappa = -0.9\sigma^2/2$ and $M >> 1$, H' drops only by about a third of its maximum value, and the mean value for δW across replications is < 0.2. So, the first set of simulations shows that with high absolute rates of mutation affecting preference and performance, populations tend to be host specialists, even when that rate is much lower than the rate of conditionally neutral mutation affecting performance, and when there are no differences in optimal performance phenotypes across hosts. Conversely, the second set of simulations show that when host environments do diverge, even when pleiotropy is more common than conditional neutrality, specialization can be prevented by reducing the absolute rate of pleiotropic mutation. With preference-and-performance pleiotropies, the preference component is key, and mutational pressure can drive arbitrary specialization across a homogeneous performance landscape.

How then do we explain the pronounced evolutionary fluctuations of host-use breadth (e.g. fig. 4)? This is not something we observed in simulations under Model 1, but the occurrence of something like these fluctuations in nature has long been suspected (Janz and Nylin 2008; Hardy 2017). Here is how we see it. Negative density-dependent viability selection pushes for broad host-use. But the pressure of mutations affecting preference-and-performance pleiotropy can undermine such generalism. Moreover, local adaptation within a host environment is more efficient when the population evolves a strong host preference (Whitlock 1996), and thus, a substantial lag of mean populations phenotypes behind host-performance optima can promote specialization. To test this intuition, we performed time-series regression analyses of 20 individual simulations with χ=1e-4, ρ=0.1, δ_O=0.005, $M=2^{-3}$, $\kappa = -0.9\sigma^2/2$, μ=1e-7. For our response variable we define an index L of the rate of adaptation that is a function of the extent to which mean performance values lagged behind host optima. Taking advantage of the fact that for each simulation there is a bimodal distribution for H', for each generation, we

classify a population as being in either a specialist or generalist state, using a cut-off value for *H'* of 0.2 When a population is in the specialist state, we define *L* as O_j - v_j*, where v_j* is the population mean. When the population is in the generalist state, *L* is ((O_1 + O_2) – (v_1* + v_2*))/2, that is, the average lag between host optima and mean performance values. As predictor variables we used *N*, *H'*, and *H*. Using the R package *metafor* (Viechtbauer and Viechtbauer, 2015) to perform a simple meta-analysis, with the estimated effects of *H'* on *L* as the effect size, and weighting effect sizes by standard errors, we found that average effect is positive and significant (effect = 0.52, p-value < 0.001); as expected, host-use generalists are slower to adapt to diverging host-use optima.

As optimal performance phenotype values in each host environment move farther away from a population's starting genotype values, it takes more alleles to produce a locally optimal phenotype (Holt et al. 2003), and the potential downside grows for recombination in the offspring of parents born in different host environments. Likewise as the proportion of pleiotropic mutation increases, it becomes harder to optimize performance in both host environments without evolving strong preference for one or the other. Thus, large displacements of O_1 and O_2 from a population's initial mean performance genotype values, along with more pervasive pleiotropy increase the time a population spends as a specialist.

Are preference-and-performance pleiotropies likely? One possibility would be a mutation at a site that regulates the expression of preference and performance genes. Another would be a mutation at a site encoding an enzyme involved in host-sensing, affecting feeding rates during development and oviposition as an adult. Yet another would be a gene for an enzyme involved in both host-sensing (e.g., degrading a host secondary compound within a sensillum on an instect's hypopharynx), and detoxification (degrades the same compound in the insect's gut). In any case, the importance of pleiotropy for the evolution of niche specialization may depend on the natural occurrence of mutations that directly affect both habitat-dependent performance and assortative mating. In that respect, our

inferences resonate with the sympatric speciation research; specialization is much more likely if there are “magic alleles” that promote local adaptation in a way that directly impedes gene flow across demes in a structured metapopulation (Bolnick and Fitzpatrick, 2007). But note that in our models the “magic alleles” that affect host-use preference as well as performance do not promote species divergence *per se*; they cause niche contraction rather than segmentation.

A few words about quantitative pleiotropy might help to contextualize our study. In classical single-locus population genetic models, alleles have direct fitness effects, and by definition an antagonistically pleiotropic allele improves fitness – or preference – for one environment and decreases it in another (Templeton 2021; e.g., Levene 1953; Rausher 1984). In contrast, in a quantitative genetics framework, alleles indirectly affect fitness via their effects on quantitative traits and the matching of trait values to environmental optima (e.g., Burger & Lynch 1995; Holt et al. 2002). Furthermore, an allele’s pleiotropic effects can be positive or antagonistic depending on the genetic environment, that is, the sum of all other allele effects across QTLs. In a word, polygenic pleiotropy is relative. Therefore, we cannot use quantitative genetics to test the predictions of the single-locus theory; that theory hinges on a concept of pleiotropy that has no direct parallel in quantitative genetics. On the other hand, since herbivorous insect host-use phenotypes are polygenic and quantitative, quantitative genetics can give us a better sense for how pleiotropy actually works in nature.

Conclusions

It would seem that pleiotropy can indeed drive quantitative niche specialization, especially when pleiotropic alleles affect preference as well as performance. On the other hand, if pleiotropies affect only performance, even with freely-evolvable habitat preference, to drive specialization they need to be much more common and antagonistic than has been observed in nature. And if habitats change quickly, or differ enough in productivity, quantitative specialization can readily evolve without pleiotropy. So, pleoiptropy is just one among the many factors that can lead to habitat specialization.

We are not the first to theoretically demonstrate that specialization can evolve without simple genetic trade-offs (e.g., Kawecki 1994; Fry 1996; Whitlock 1996; Kawecki et al. 1997; Dragi 2021). But assuming monogeny, the joint evolution of niche preference and performance has seemed an especially powerful mechanism (Ravigne et al. 2007). Our models show under more realistic polygenic architectures, the influence of pleiotropy is more limited. And so we suggest that future theoretical and experimental work might look deeper into that extent of pleiotropy on habitat preference and performance, and more broadly for what stands in the way of the evolution of generalism (Hardy et al., 2020).

References

Bolnick, D. I., & Fitzpatrick, B. M. (2007). Sympatric speciation: models and empirical evidence. *Annu. Rev. Ecol. Evol. Syst.*, 38, 459–487. https://doi.org/10.1146/annurev.ecolsys.38.091206.095804

Burger, R., & Lynch, M. (1995). Evolution and extinction in a changing environment: a quantitative genetic analysis. Evolution, 49(1), 151–163.

Draghi, J. A. (2021). Asymmetric evolvability leads to specialization without trade-offs. The American Naturalist, 197(6), 000-000. https://doi.org/10.1086/713913

Egan, S. P., Ragland, G. J., Assour, L., Powell, T. H. Q., Hood, G. R., Emrich, S., Nosil, P., & Feder, J. L. (2015). Experimental evidence of genome-wide impact of ecological selection during early stages of speciation-with-gene-flow. *Ecology Letters*, *18*(8), 817–825. https://doi.org/10.1111/ele.12460

Egas, M., Dieckmann, U., & Sabelis, M. W. (2004). Evolution restricts the coexistence of specialists and generalists: The role of trade-off structure. *American Naturalist*, *163*(4), 518–531. https://doi.org/10.1086/382599

Felsenstein, J. (1981). Skepticism towards Santa Rosalia, or why are there so few kinds of animals? *Evolution* 35:124–138. https://doi.org/10.2307/2407946

Forister, M. L., Novotny, V., Panorska, A. K., Baje, L., Basset, Y., Butterill, P. T., Cizek, L., Coley, P. D., Dem, F., Diniz, I. R., Drozd, P., Fox, M., Glassmire, A. E., Hazen, R., Hrcek, J., Jahner, J. P., Kaman, O., Kozubowski, T. J., Kursar, T. A., … Dyer, L. A. (2015). The global distribution of diet

breadth in insect herbivores. *Proceedings of the National Academy of Sciences of the United States of America*, 112(2), 442–447. https://doi.org/10.1073/pnas.1423042112

Fry, J. D. (1996). The Evolution of Host Specialization: Are Trade-Offs Overrated? *The American Naturalist*, 148, S84–S107. https://doi.org/10.1086/285904

Futuyma, D. J., & Moreno, G. (1988). The Evolution of Ecological Specialization. *Annual Review of Ecology and Systematics*, 19, 207–233.

Gompert, Z., Jahner, J. P., Scholl, C. F., Wilson, J. S., Lucas, L. K., Soria-Carrasco, V., Fordyce, J. A., Nice, C. C., Buerkle, C. A., & Forister, M. L. (2015). The evolution of novel host use is unlikely to be constrained by trade-offs or a lack of genetic variation. *Molecular Ecology*, *24*(11), 2777–2793. https://doi.org/10.1111/mec.13199

Haller, B. C., & Messer, P. W. (2019). SLiM 3: Forward Genetic Simulations Beyond the Wright-Fisher Model. *Molecular Biology and Evolution*, 36(3), 632–637. https://doi.org/10.1093/molbev/msy228

Hardy, N. B. (2017). Do plant-eating insect lineages pass through phases of host-use generalism during speciation and host switching? Phylogenetic evidence. *Evolution*, 71(8), 2100-2109. https://doi.org/10.1111/evo.13292

Hardy, N. B., Kaczvinsky, C., Bird, G., & Normark, B. B. (2020). What We Don't Know About Diet-Breadth Evolution in Herbivorous Insects. *Annual Review of Ecology, Evolution, and Systematics*, 51(1). https://doi.org/10.1146/annurev-ecolsys-011720-023322

Holt, R.D., Gomulkiewicz, R., & Barfield, M. (2003). The phenomenology of niche evolution via quantitative traits in a 'black-hole'sink. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(1511), 215–224.

Jaenike, J. (1990). Host Specialization in Phytophagous Insects. *Annual Review of Ecology and Systematics*, 21, 243–273.

Levene, H. (1953). Genetic equilibrium when more than one niche is available. *The American Naturalist*, 87, 331–333.

Liu, H., Jia, Y, Sun, X., Tian, D., Hurst, L. D., & Yang, S. (2016) Direct Determination of the mutational rate in bumblebee reveals evidence for weak recombination-associated mutation and an approximate rate constancy in insects. *Molecular Biology and Evolution*, 34(1), 119–130.

Lynch, M. (2011). Statistical inference of the mechanisms of genome evolution. *PLoS Genetics*, 7(6), e1001389. https://doi:10.1371/journal.pgen.1001389

Nylin, S., & Janz, N. (2008). 15. The Oscillation Hypothesis Of Host-Plant Range And Speciation. In *Specialization, Speciation, and Radiation* (pp. 203-215). University of California Press. https://doi.org/10.1525/9780520933828-017

Kawecki, T. J. (1994). Accumulation of deleterious mutations and the evolutionary cost of being a generalist. *The American Naturalist*, 144(5), 833–838. https://doi.org/10.1086/285709

Kawecki, T.J., Barton, N., & Fry, J. (1997). Mutational collapse of fitness in marginal habitats and the evolution of ecological specialisation. *Journal of Evolutionary Biology*, 10(3), 407–429.

Kirkpatrick, M., & Ravigné, V. (2002). Speciation by natural and sexual selection: models and experiments. *The American Naturalist*, 159(S3), S22-S35. https://doi.org/10.1086/338370

Rainey, P. B., Buckling, A., Kassen, R., & Travisano, M. (2000). The emergence and maintenance of diversity: insights from experimental bacterial populations. *Trends in Ecology & Evolution*, 15(6), 243–247. https://doi.org/10.1016/S0169-5347(00)01871-1

Rausher, M. D. (1984). The evolution of habitat preference in subdivided populations. *Evolution*, 38, 596–608.

Ravigné, V., Dieckmann, U., & Olivieri, I. (2009). Live where you thrive: joint evolution of habitat choice and local adaptation facilitates specialization and promotes diversity. *American Naturalist*, 174(4). https://doi.org/10.1086/605369

Rueffler, C., Van Dooren, T. J. M., & Metz, J. A. J. (2004). Adaptive walks on changing landscapes: Levins' approach extended. *Theoretical Population Biology*, 65(2), 165–178. https://doi.org/10.1016/j.tpb.2003.10.001

Smith, J. M., & Hoekstra, R. (1980). Polymorphism in a varied environment: how robust are the models? *Genetics Research*, 35(01), 45. https://doi.org/10.1017/S0016672300013926

Savolainen, O., Lascoux, M., & Merila, J. (2013). Ecological genomics of local adaptation. *Nature Reviews Genetics*, 14, 807–820. https://doi.org/10.1038/nrg3522

Templeton, A. (2021). Population Genetics and Microevolutionary Theory. Hoboken, New Jersey: John Wiley &Sons.

van Tienderen, P. H. (1991). Evolution of generalists and specialists in spatially heterogeneous environments. *Evolution*, 45:1317–1331.

van Tienderen, P. H. (1997). Generalists, specialists, and the evolution of phenotypic plasticity in sympatric populations of distinct species. *Evolution* 51:1372–1380.

Vertacnik, K. L., & Linnen, C. R. (2017). Evolutionary genetics of host shifts in herbivorous insects: insights from the age of genomics. *Annals of the New York Academy of Sciences*, 1389(1), 186–212. https://doi.org/10.1111/nyas.13311

Whitlock, M. C. (1996). The red queen beats the jack-of-all-trades: the limitations on the evolution of phenotypic plasticity and niche breadth. *The American Naturalist*, 148, S65-S77. https://doi.org/10.1086/285902

Wilfert, L., Gadau, J., & Schmid-Hempel, P. (2007). Variation in genomic recombination rates among animal taxa and the case of social insects. *Heredity*, 98(4), 189–197. https://doi.org/10.1038/sj.hdy.6800950

Xu, L., Chen, H., Hu, X., Zhang R., & Luo, Z. W. (2006). Average gene length is highly conserved in prokaryotes and eukaryotes and divergences only between the two kingdoms. *Molecular Biology and Evolution*, 23(6):1107–1108.

Table 1. Parameter definitions and values; variable definitions and intervals.

Parameters

Symbol	**Definition**	**Values**
μ	Mutation rate	{1e-7; 1e-9}
χ	Recombination rate	{1e-4; 1e-6}
λ	Mean and variance of Poisson-distributed fecundity	{1.5}
n	Number of host species	{2}
σ^2	Allele effect variance	{0.6}
κ	Pleiotropic effect covariances	$\{-0.9\sigma^2, 0, 0.9\sigma^2\}$
M	Mix of pleiotropic to conditionally neutral mutations	$\{0, 2^{-3}, 2^0, 2^3, 2^6\}$
ρ	Mix of preference to performance mutations	{0, 0.1, 1}
ω	Gaussian fitness function variance	{3.0}
δ_O	Per-generation rate of change in performance optima	{0.005, 0.05, 0.1}
K_j	Carrying capacity of host species j	{500}
pK_1	Proportional carrying capacity of host 1	{0.05, 0.25, 0.5}

Variables

Symbol	**Definition**	**Intervals**
h_j	Host species j	$\{h_1, h_2\}$
N	Total population size	$0 \leq N$, with a soft maximum of ΣK_j
N_j	Population size on host species j	$0 \leq N_j$, with a soft maximum of K_j
z_{ij}	Performance genotype value of individual i on h_j	$z_{ij} \in \mathbb{R}$
y_1	Genotype value for preference for h_1	$y_1 \in \mathbb{R}$
O_j	Performance optimum for h_j	$0 \leq O_j \leq \delta_O$ * post-burnin generations
H	Population heterozygosity	$0 \leq H \leq 1$
H'	Shannon's diversity of host-use	$0 \leq H' \leq \ln(n)$
$v(z_{ij})$	Genotype-environment match viability	$0 \leq v(z_{ij}) \leq 1$
W	Mean population fitness across host species	$0 \leq W \leq 1$
F_{ST}	Index of neutral genetic subpopulation divergence	$0 \leq F_{ST} \leq 1$
δP	Difference in mean y_i between subpopulations	$0 \leq \delta P$
δW	Specialization index	$0 \leq \delta_W \leq 1$

Figures

Figure 1. Simulated herbivore life cycle. Each generation consists of (1) a round of random mating followed by (2) offspring production, in which parent genomes crossed and mutated, and then offspring are added to the parental host subpopulation. After reproducing, parents die. The environment consists of two host plant species. Offspring chose to stay on their natal host species, or (3) migrate to the other host species. After a host is chosen, (4) viability is first determined by the match between an individual's performance phenotype and their host environment, and then scaled by local subpopulation density.

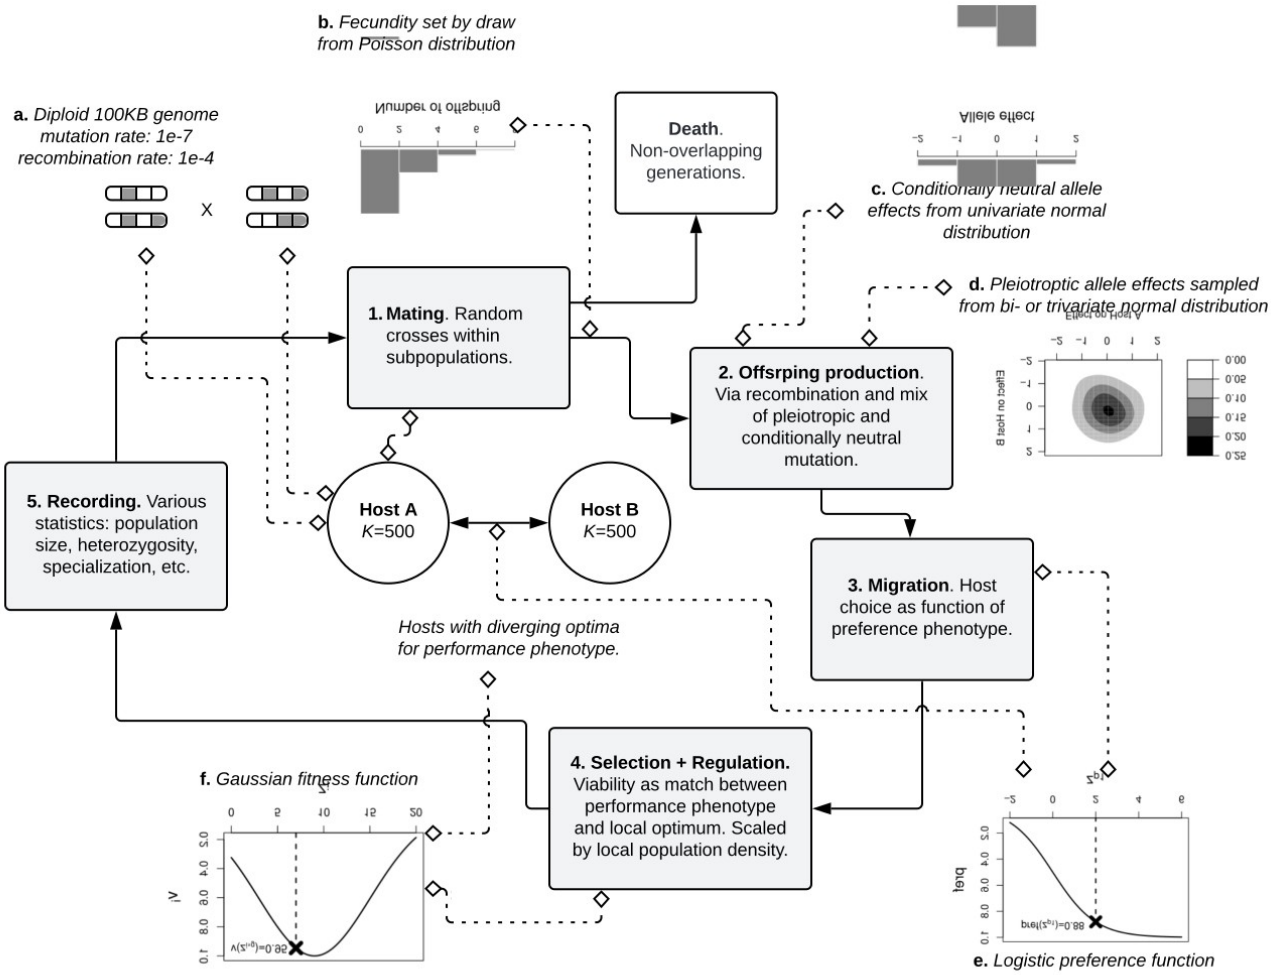
a. Diploid 100KB genome
mutation rate: 1e-7
recombination rate: 1e-4
X
b. Fecundity set by draw
from Poisson distribution
Number of offspring
Death. Non-overlapping generations.
Allele effect
c. Conditionally neutral allele effects from univariate normal distribution
d. Pleiotroptic allele effects sampled from bi- or trivariate normal distribution
Effect on host A
Effect on host B
1. Mating. Random crosses within subpopulations.
2. Offsrping production. Via recombination and mix of pleiotropic and conditionally neutral mutation.
Host A
K=500
Host B
K=500
5. Recording. Various statistics: population size, heterozygosity, specialization, etc.
3. Migration. Host choice as function of preference phenotype.
Hosts with diverging optima
for performance phenotype.
4. Selection + Regulation. Viability as match between performance phenotype and local optimum. Scaled by local population density.
f. Gaussian fitness function
e. Logistic preference function

Figure 2. Example runs under Model 1. The left column (a–d) shows a simulation in which host species are equally productive (K_1=K_2), preference mutations are as common as performance mutations (ρ=1), and in which optimal performance values change slowly (δ_O=0.005). The middle column (e–h) shows a simulation in which K_1=K_2, ρ=1, and optimal performance value change rapidly (δ_O=0.1). The right column (i–l) shows a simulation in which K_1<K_2 and optimal performance values change slowly. The top row shows changes in population mean values for performance traits (z_1 and z_2), over 1000 generations after the burn-in period. The color of the tracing line changes with a population's mean host-performance specificity; lighter shades of gray denote higher specificity. The dashed orange arrow shows the concomitant change in host performance optima (O_j). The green 'X' denotes the population mean at the end of the simulation.

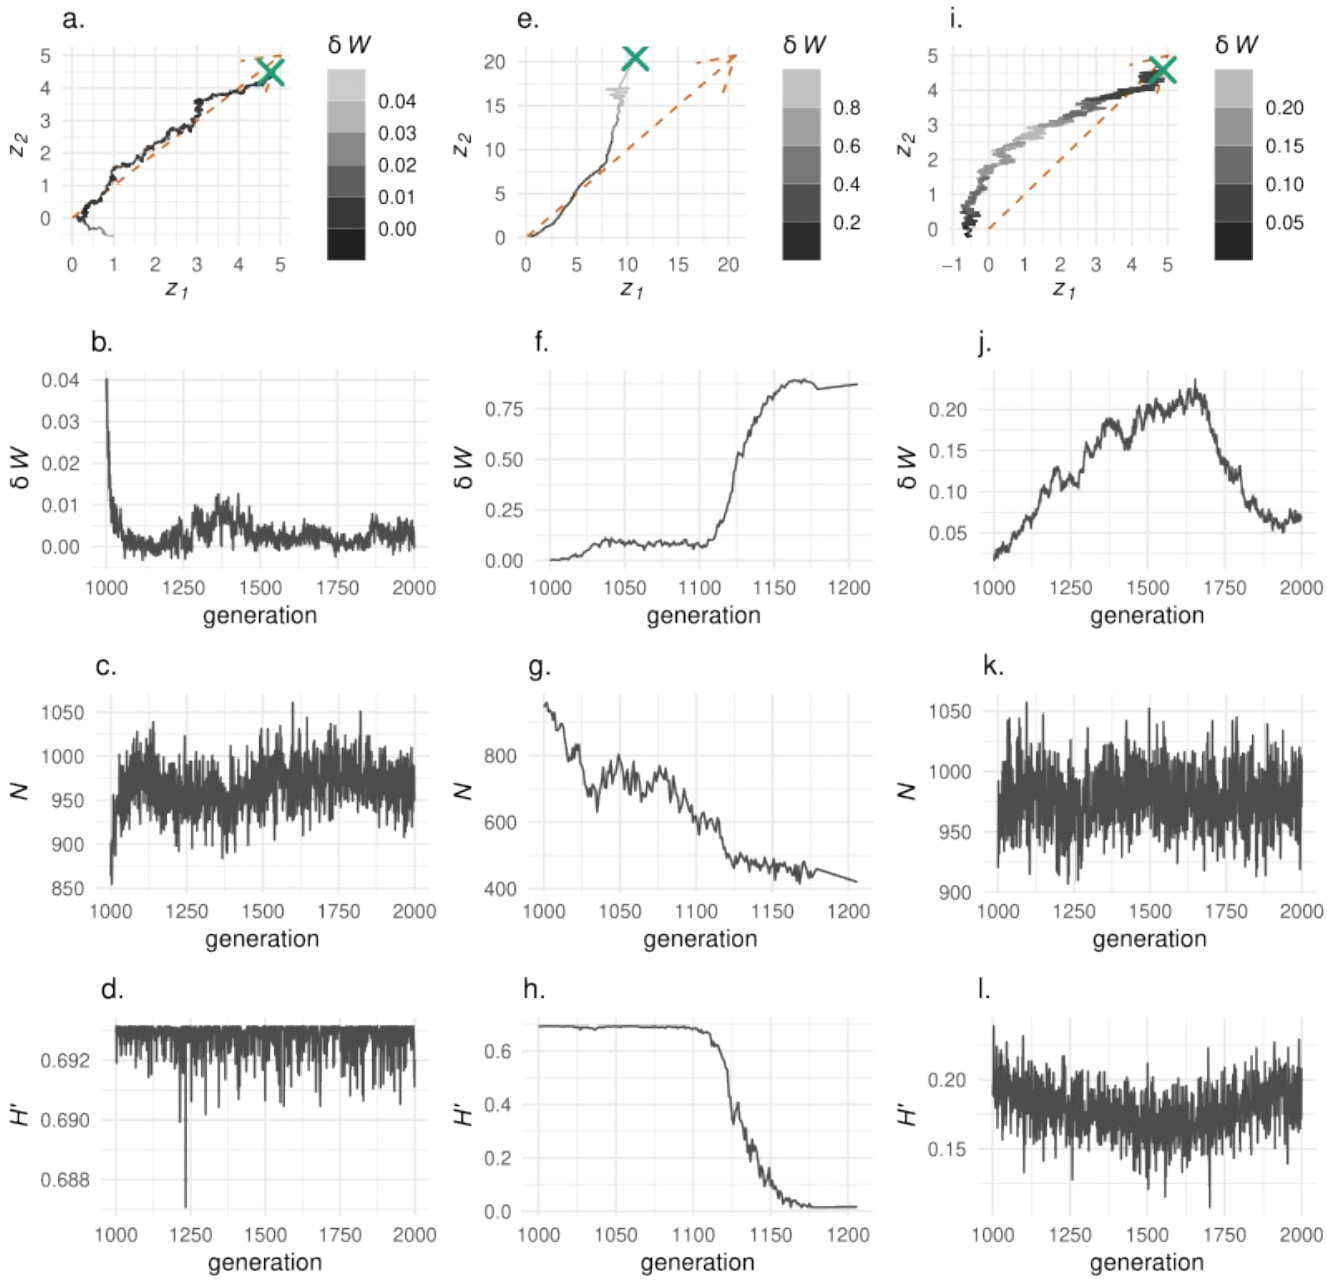
a.
e.
i.
δ W
z2
z1
b.
f.
j.
c.
g.
k.
d.
h.
l.
N
H'
generation

Figure 3. Summary of runs under Model 1 with ρ=1. *M* is the ratio of pleiotropic to conditionally neutral mutation. Green lines are for pleiotropic covariances of -0.54, orange lines are for covariances of zero, and purple lines are for covariances of 0.54. All plots show mean population states over 1000 generations after the burn-in period. The left column (a–d) shows population states when host species are equally productive (K_1=K_2) preference mutations are as common as performance mutations (ρ=1), and in which optimal performance values change slowly (δ_O=0.005). The middle column (e–h) is for simulations in which K_1=K_2, ρ=1, and optimal performance value change rapidly (δ_O=0.1). The right column (i–l) is for simulations in which K_1<K_2 (pK_1 = 0.25) and optimal performance values change slowly.

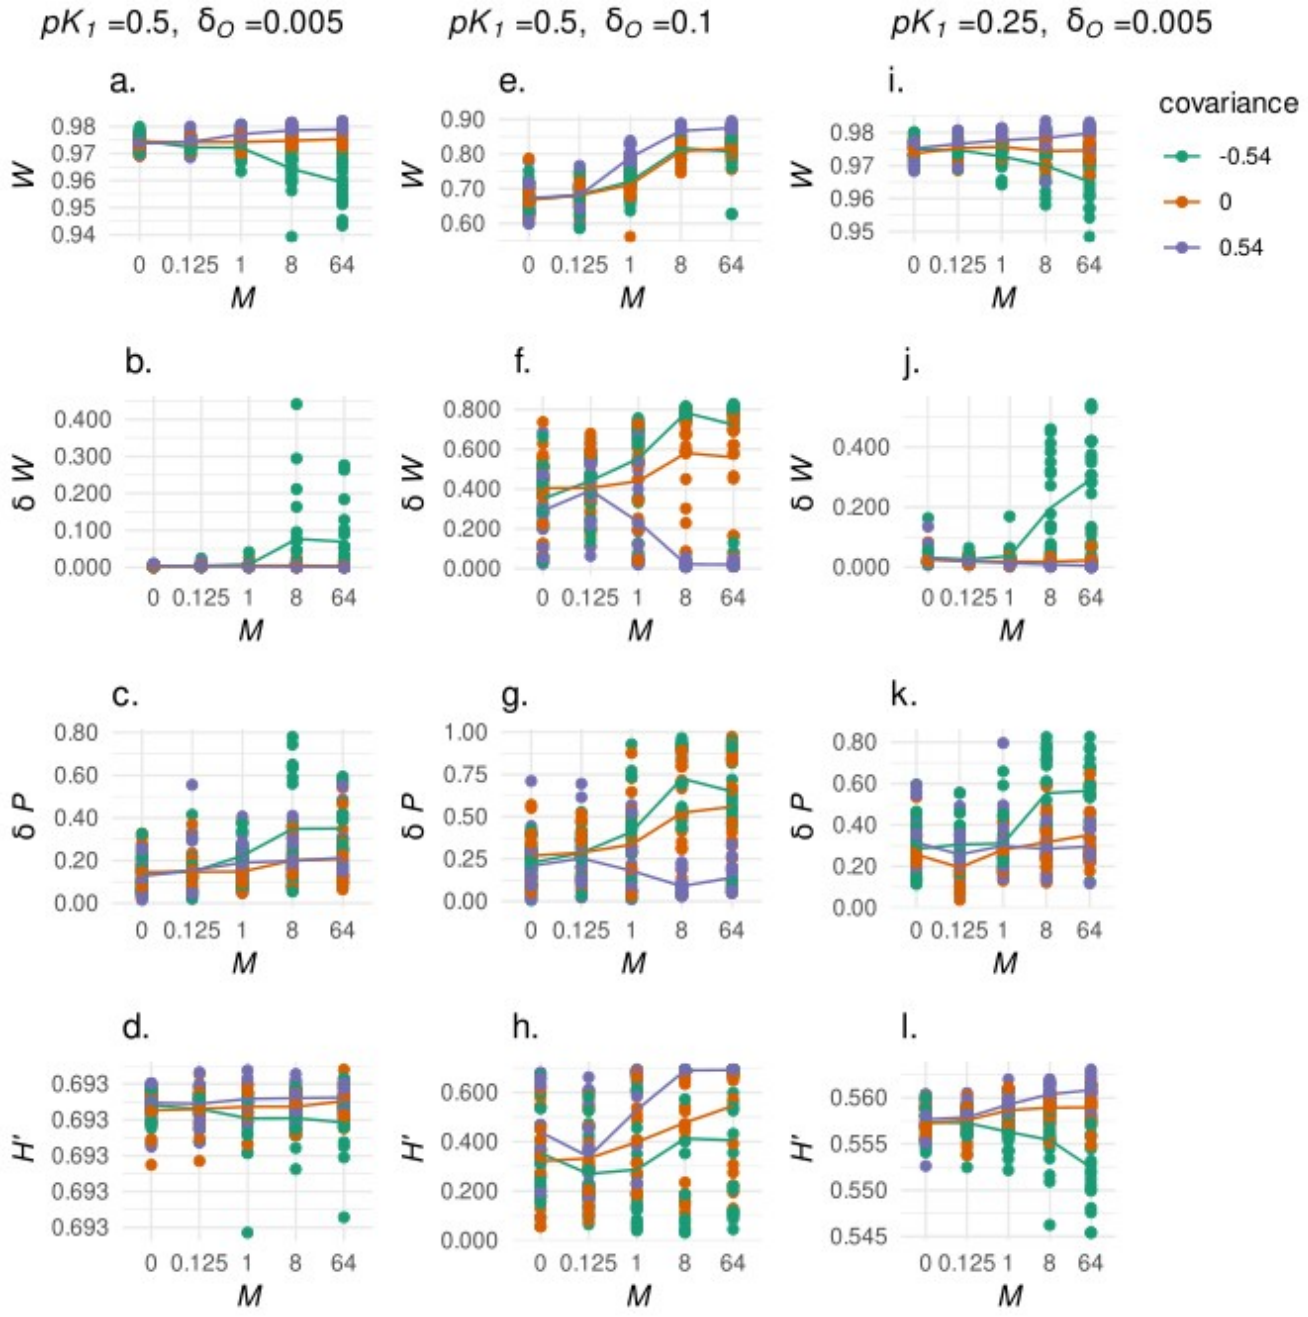
pK1 =0.5, δO =0.005
pK1 =0.5, δO =0.1
pK1 =0.25, δO =0.005
a.
b.
c.
d.
e.
f.
g.
h.
i.
j.
k.
l.
W
δ W
δ P
H'
M
0 0.125 1 8 64
covariance
-0.54
0
0.54

Figure 4. Example run under Model 2. Values for model parameters are χ = 1e-4, ρ = 1, κ = - 0.9σ^2/2, $M = 2^{-3}$, and δ_O = 0.005. (a) Changes in population mean values for performance traits (z_1 and z_2), over 1000 generations after the burn-in period. The color of the tracing line changes with a population's mean host-performance specificity; lighter shades of gray denote higher specificity. The dashed orange arrow shows the concomitant change in host performance optima (O_j). The green 'X' denotes the population mean at the end of the simulation. Note the pronounced deviations of the populaiton's evolutionary path from the optimal vector. The other three panels show the evolution of δW (b), N (c), and H' (d) following the burn-in period. In (d) the color of the tracing line changes with the proportion of the population found on h_1; this allows us to see that the population temporarily specializes on different hosts.

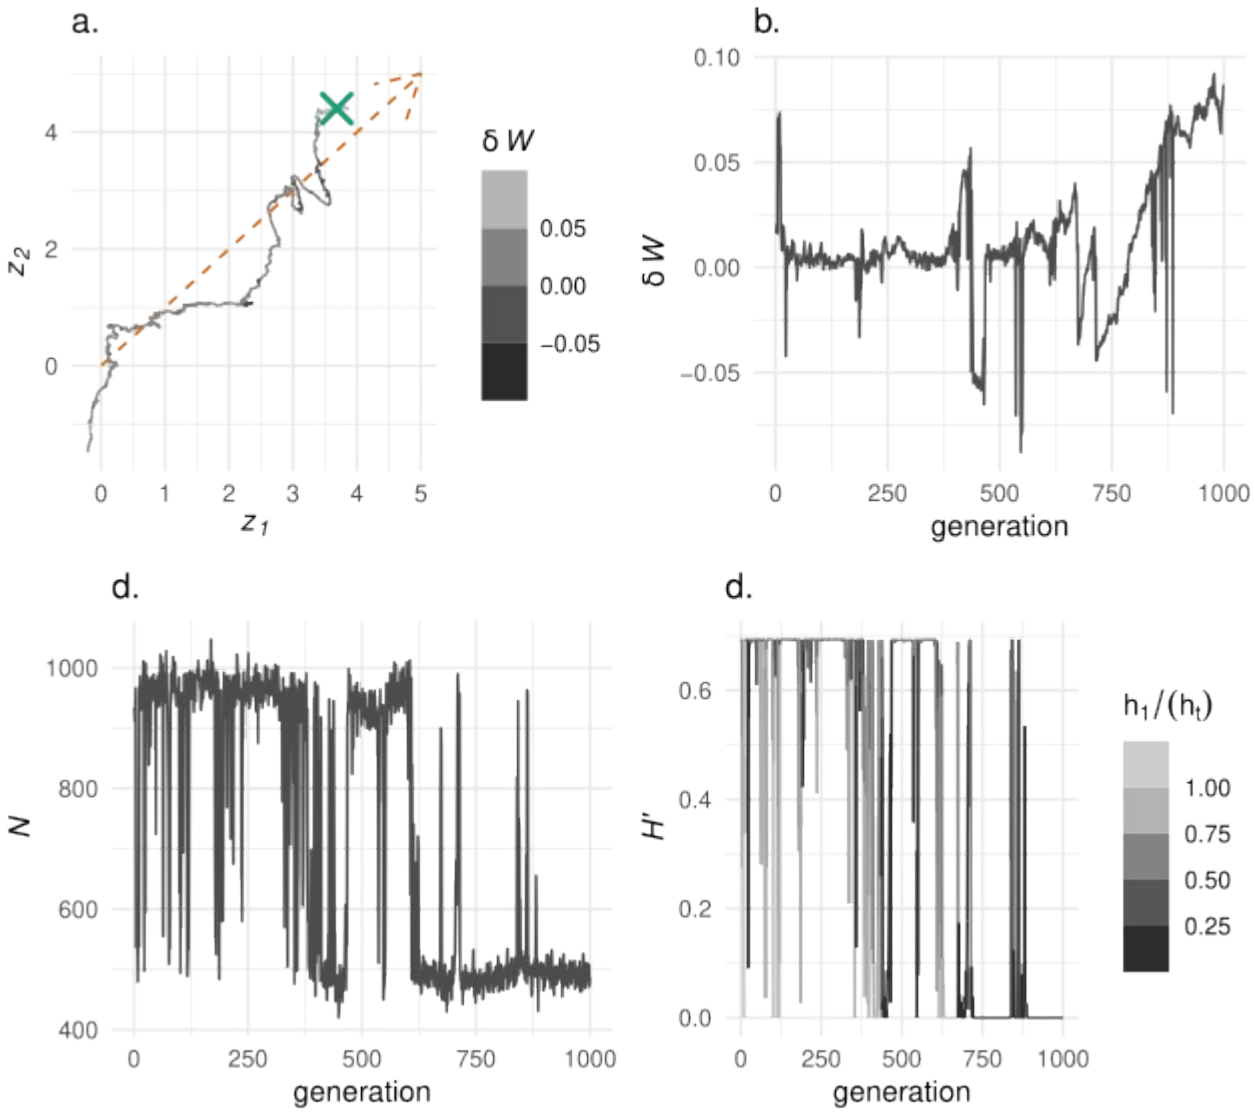
a.
b.
d.
d.
δ W
0.05
0.00
−0.05
z2
z1
δ W
generation
N
H'
h1/(ht)
1.00
0.75
0.50
0.25

Figure 5. Summary of runs under Model 2. *M* is the ratio of pleiotropic to conditionally neutral mutation. Green lines are for pleiotropic covariances of -0.54, orange lines are for covariances of zero, and purple lines are for covariances of 0.54. All plots show mean population states over 1000 generations after the burn-in period. The top row (a–d) shows population states when χ = 1e-4, ρ = 1, μ = 1e-7 and δ_O = 0, that is, host performance optima are fixed at the same value for both host species. The bottom row (e–h) is for simulations in which χ = 1e-4, ρ = 1, μ = 1e-9 δ_O = 0.005, that is with slowly changing host performance optima, but reduced rates of mutation.

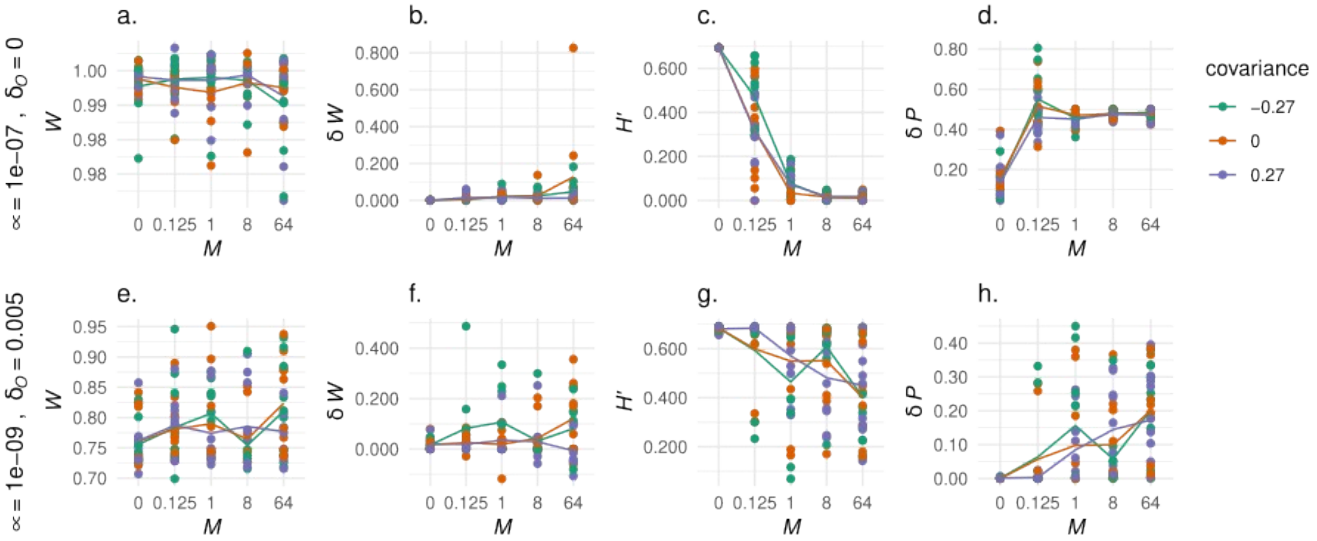
a.
b.
c.
d.
e.
f.
g.
h.
∝ = 1e−07 , δO = 0
∝ = 1e−09 , δO = 0.005
W
δ W
H'
δ P
M
0 0.125 1 8 64
covariance
−0.27
0
0.27

Fig S1. Pleiotropic allele effects on performance. In Model 1, pleiotropic allele effects on host-use performance are drawn from a bivariate normal distribution. Here we show 100 random deviates from a distribution with means of zero, variances of 0.6, and covariences of 0.54. This is one of the three distributions used in our simulations. General vigor alleles are shown with a "+", General weakness alleles with a "o", and antagonistically plieotropic alleles with an "x." In this case, most of the sampled alleles are antagonistically pleiotropic. At the start of each simulation, the optimal performance phenotype value in each host environment is zero. Then, after the burn-in period O_j increases by δ_O each generation.

Figure S2. Summary of runs under Model 1, with ρ=0.1. Green lines are for pleiotropic covariances of -0.54, orange lines are for covariances of zero, and purple lines are for covariances of 0.54. All plots show mean population states over 1000 generations after the burn-in period. The left column (a–d) shows population states when host species are equally productive (K_1=K_2), and in which optimal performance values change slowly (δ_O=0.005). The middle column (e–h) is for simulations in which K_1=K_2, and optimal performance value change rapidly (δ_O=0.1). The right column (i–l) is for simulations in which K_1<K_2 (pK_1 = 0.25) and optimal performance values change slowly.

Figure S3. Summary of runs under Model 1, with χ=1e-6. Green lines are for pleiotropic covariances of -0.54, orange lines are for covariances of zero, and purple lines are for covariances of 0.54. All plots show mean population states over 1000 generations after the burn-in period. The left column (a–d) shows population states when host species are equally productive (K_1=K_2), and in which optimal performance values change slowly (δ_O=0.005). The middle column (e–h) is for simulations in which K_1=K_2, and optimal performance value change rapidly (δ_O=0.1). The right column (i–l) is for simulations in which K_1<K_2 (pK_1 = 0.25) and optimal performance values change slowly.